HIPPI / Serial-HIPPI

Don E. Tolmie
Los Alamos National Laboratory, Los Alamos, New Mexico, (505) 667-5502

Marty G. Halvorson
Digital Equipment Corporation, Los Alamos, New Mexico, (505) 662-2011

Abstract

The High-Performance Parallel Interface (HIPPI) is a simple high-performance, point-to-point channel for transmitting digital data at peak data rates of 800 or 1600 Mbit/s. The transmission distance between data-processing equipment using copper cabling can be up to 25 meters. This distance may be increased by using a Serial-HIPPI extender. This paper describes the HIPPI channel, the choices considered for a serial extender, and the resulting agreement for the Serial-HIPPI specification.

1: Introduction

The High-Performance Parallel Interface - Mechanical, Electrical, and Signalling Protocol (HIPPI-PH) is American National Standard X3.183-1991 [1]. It specifies the physical layer of an 800 or 1600 Mbit/s (100 or 200 MByte/s) point-to-point channel for interconnecting computers and other data-processing equipment. HIPPI-PH specifies a copper twisted-pair cable for distances up to 25 meters. The signalling protocol, including the flow control, is useful for distances of many kilometers. The Serial-HIPPI specification came from the desire of the participants in the ANSI X3T9.3 Task Group developing the HIPPI standards to provide a single defined way to extend HIPPI-PH channel distances beyond 25 meters by using fiber optic components.

Other HIPPI documents include (1) HIPPI-FP (Framing Protocol) - much like a data link, (2) HIPPI-LE (Link Encapsulation) - a mapping to IEEE 802.2 so that network protocols such as TCP/IP can operate over HIPPI, (3) HIPPI-MI (Memory Interface) - with read, write, and test and set operations for accessing a memory over HIPPI, and (4) HIPPI-SC (Switch Control) - control of circuit, e.g., crossbar, switches interconnecting multiple HIPPI-PH channels. In addition, operation of the IPI-3 command sets for disks and tapes over HIPPI are also being specified within the IPI-3 standards documents. The Serial-HIPPI operates at the HIPPI-PH level without affecting any of the other HIPPI functions except for the addition of small delays incurred in the longer distances possible with fiber optics. In this paper the HIPPI abbreviation is used to mean HIPPI-PH.

Many vendors are providing HIPPI channels on their equipments today. HIPPI is currently the interface of choice for high-end applications.

2: HIPPI Characteristics

HIPPI was designed to move data from memory to memory at such high rates that it has been likened to a fire hose. An early guideline was to keep the design simple. As such, it looks more like a communications channel than like many traditional computer channels.

In relation to the Open Systems Interconnection (OSI) Basic Reference Model, HIPPI-PH covers the physical layer. HIPPI-PH uses a parallel data path with copper twisted-pair cable. The 800-Mbit/s HIPPI uses a 32-bit data bus and one cable. The 1600-Mbit/s version uses a 64-bit data bus and two cables. The major emphasis has been on the development of the 800-Mbit/s version.

HIPPI is a simplex channel, capable of transferring data in one direction only. Two HIPPIs may be used to implement a full-duplex channel. HIPPI is a point-to-point channel that does not support multi-drop. The point-to-point limitation considerably simplified the electrical and protocol aspects of the HIPPI. Crossbar switches and other networking methods may be used to achieve the equivalent of multi-drop. An addressing mechanism is included in HIPPI to support these networking concepts.

The HIPPI signal sequences provide look-ahead flow control to allow the average data rate to approach the peak data rate, even over distances of tens of kilometers. Data transfers and flow control are performed in increments of

0-8186-2655-0/92 $3.00 © 1992 IEEE

bursts, with each burst nominally containing 256 words (1024 or 2048 bytes).

Error detection, but not error correction, is provided by HIPPI. Byte parity is used on the data bus. In addition, each burst of data, where a burst is 256 words or less, is immediately followed by a length/longitudinal redundancy checkword (LLRC). Any error recovery is done by higher-layer protocols. Driving a video frame buffer, for example, is best served by ignoring errors that the next frame overwrites anyway.

HIPPI also provides support for low-latency, real-time, and variable-size packet transfers. The signal-line control sequences are simple and do not require any custom silicon to implement.

3: HIPPI Data Framing and Signalling Sequences

Figure 1 shows the basic organization of the information or data framing on HIPPI. A connection is made in a fashion similar to the connection made when dialing the telephone. Once a connection is established a packet (or multiple packets) can be sent from the source to the destination. Each packet contains one or more bursts, and each burst contains one to 256 words. Bursts that contain less than 256 words may only occur as the first or last burst of a packet. Words are composed of 32 or 64 bits. The amount of wait time between packets and bursts may vary. Maximum wait times depend on the data flow to or from the upper-layer protocols and on the data flow to or from the opposite end of the channel.

The interface signals are illustrated in Figure 2. The numbers in parentheses indicate the number of signal lines when using the 1600-Mbit/s option. The other numbers indicate the number of signal lines when using the 800-Mbit/s option. All signals, except for the INTERCONNECT signals, use differential emitter-coupled logic (ECL) drivers and receivers. The INTERCONNECT signals use single-ended ECL drivers and receivers.

Fifty-pair, twisted-pair cables are specified for distances up to 25 meters. 100-pin thin-line connectors are used. All of the signal lines in the HIPPI are unidirectional. The control and data signals are timed in relation to the constant 25-MHz CLOCK signal with a period of 40 nanoseconds.

Typical HIPPI waveforms are shown in Figure 3 for a sequence that establishes a connection, sends a packet containing two bursts, sends a packet containing one burst, and then disconnects. A connection is made from the source to the destination much like a telephone connection. The source supplies the I-Field on the data bus (like a telephone number), and asserts the REQUEST signal. If the destination wants to accept the connection, it asserts the CONNECT signal. The format of the I-Field for controlling physical layer switches is specified in the HIPPI-SC document.

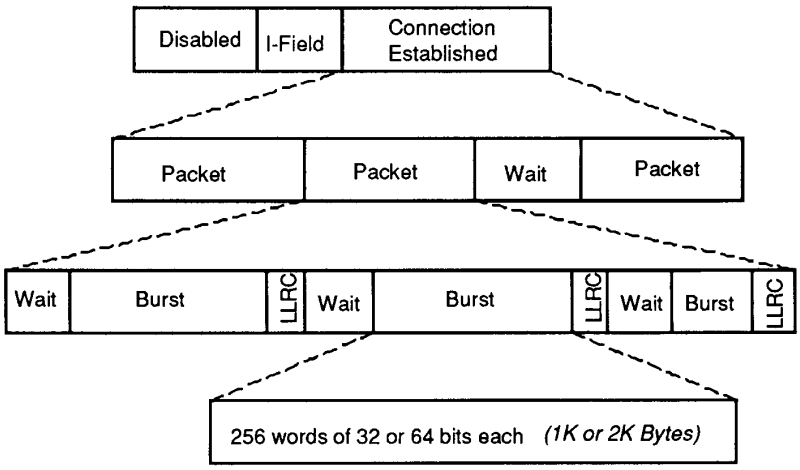

Figure 1. HIPPI framing hierarchy

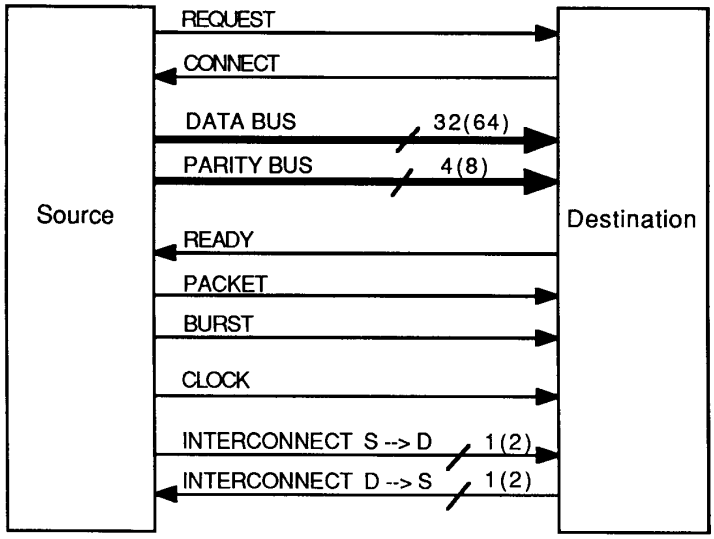

The numbers indicate the number of signal lines when using the 800-Mbit/s (1600-Mbit/s) option.

Figure 2. Summary of the HIPPI interface signals.

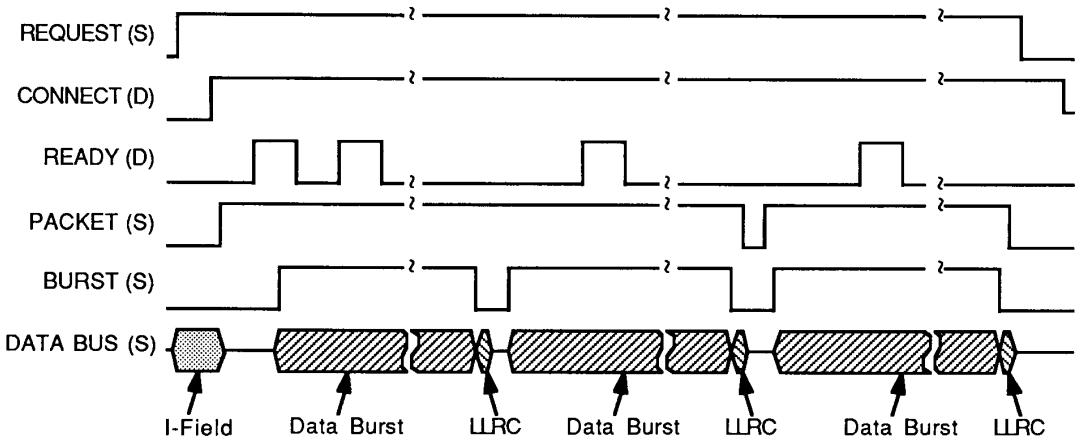

Figure 3. Typical HIPPI waveforms.

Once a connection is established, single or multiple packets may be transferred from the source to the destination. Packets are delimited by the PACKET signal being true.

Bursts are delimited by the BURST signal being true. Bursts consist of a group of one to 256 words sent on the data bus, one word per clock period, during contiguous clock periods. The LLRC checkword is sent from the source to the destination on the data bus during the first clock period following the burst.

The destination controls the flow of data by issuing a READY pulse for each burst that it is prepared to accept from the source. If the READY pulses arrive at the source before the source is ready to send the next burst, there will be no time lost between bursts. Hence, the flow control is distance independent if the cable-length time is shorter than the time required to transmit the number of bursts. This requires about one burst buffer for each two kilometers of cable distance.

The INTERCONNECT signals indicate to both the source and destination that the cable(s) are connected and that the other end is powered up. They may also be used to indicate whether the 800-Mbit/s or 1600-Mbit/s HIPPI option is in use.

4: Serial-HIPPI overview and history

4.1: Overview

The Serial-HIPPI [2] as shown in Figure 4, provides a full-duplex path between dual HIPPI channels in physically separated hosts. Figure 4 shows the Serial-HIPPI interfaces in an outboard configuration, connected to the hosts by 50-pair copper cables up to 25 meters long. Serial-HIPPI can also be built into the host, omitting the copper cables and their connectors.

The function of the Serial-HIPPI is to serialize the HIPPI channel for transmission over a pair of fibers or coax cables. The objectives of Serial-HIPPI are:

- Compatibility: Serial-HIPPI can be treated as virtual ribbon cable for HIPPI-PH. Other than the extended range, the user should see no difference in operation when using Serial-HIPPI.

- Interoperability: Systems from various vendors that meet the Serial-HIPPI specification should be plug compatible at the serial electrical, serial optical, or parallel HIPPI interfaces.

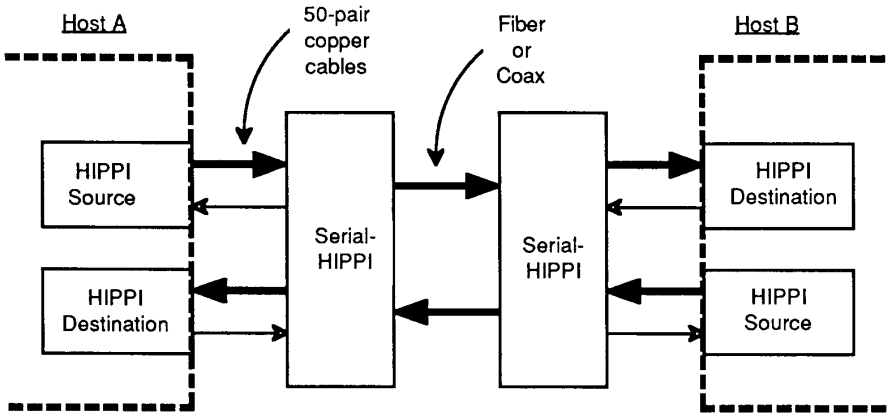

Figure 4. Outboard Serial-HIPPI

4.2: Need

Serial-HIPPI began in response to a need for HIPPI data rates over distances exceeding the 25 meters available with parallel HIPPI. Examples of these needs are:

- Scientific visualization outside the confines of the noisy, disruptive computer room environment containing the machines powerful enough to compute the visualization frames.

- Data acquisition where the test sites are remote from the processing computers. For example, the CERN experiment sites are located approximately 10 km from the central computing facility [3]. These sites generate on the order of 2 terabyte/s. After hardware processing, approximately 2 gigabyte/s are (will be) sent to the central site for further processing.

- Distributed algorithms are being developed to take advantage of different computer architectures to process that part of a problem most effectively solved by each different architecture. For example, the CASA testbed is developing an algorithm for seismic modeling that distributes the problem among scalar super-computers and massively parallel super-computers for simultaneous solution [4,5].

High-end workstations are capable of HIPPI data rates but are, generally, severely limited in connector space. Thus there is also a need for a copper, coax based solution.

4.3: History and ANSI affiliation

The call for participants in developing Serial-HIPPI went out in May 1990 to members of X3T9.3. Sufficient interest was expressed so that the first meeting took place in October 1990. One of the authors (Halvorson) agreed to serve as chairman of the group.

Another ANSI effort, Fibre Channel, was also in progress in X3T9.3, and it would also provide a solution. However, that effort wasn't expected to be complete until sometime in 1992 at the earliest.

The Serial-HIPPI group felt that to satisfy the immediate need, a interim agreement, reached in a minimum time period, utilizing HIPPI-PH as the input and output media would provide the best near term solution. Further, since ANSI affiliation implies certain administrative requirements and their attendant overhead, it was felt that ANSI affiliation would impede progress toward a quick solution. For these reasons, the stated goal of the Serial-HIPPI group was to develop an agreement on a specification that:

- hardware vendors could implement using existing technology,
- would provide inter-operability at the lowest level of hardware,
- would be completed in a minimum amount of time, and
- would not be an ANSI standard.

Based on the goals, a charter was written that read:

> "Define a serial I/O specification at the fiber and coax level that offers transparent transport of HIPPI protocols, and is inter-operable between different vendors."

The completed agreement was finished in April 1991, a total of 11 months from call for participants to publication of the specification. For comparison, the ANSI effort resulting in HIPPI-PH took approximately 4 years.

5: Solutions

There were a number of solutions to providing inter-operable Serial-HIPPI proposed. From these proposals one solution was accepted.

5.1: Requirements

The Serial-HIPPI group established a set of system parameters that any proposed solution must meet. These parameters are:

- Distance - 10 km on fiber, 25 m on coax
- Bit error rate to user - less than 10^{-12}
- Dual simplex operation is normal mode. Support for one way HIPPI is required (i.e. simplex).
- 1600 Mbit/s HIPPI support is required
- Data rate - 1 Gbit/s which includes HIPPI data, parity, and control signals
- Fiber plant - single mode (EIA 492-BAAA) which is the same as selected by Fibre Channel
- Maximum skew among fibers in same bundle - 22 ns. Used in 1600 Mbit/s operation

The group also established a set of selection criteria for the encoding scheme. These criteria included:

High priority
- High availability (goal was April 1991)
- User bit error rate (BER) less than 10^{-12}
- Parts be second sourced

Medium priority
- Reasonable system price (reasonable was not defined)
- Circuit board real estate requirements
- Power consumption

Low priority
- Ability to drive copper (i.e. coaxial cable) links
- Resolution of issues regarding patents, licensing, etc.

5.2: Proposed solutions

The proposed solutions for Serial-HIPPI were:

- Gazelle Microcircuits proposed a solution using their "Hot Rod" © chip. This solution made use of the 4B/5B encoding used in FDDI.
- Ancor Communications proposed a solution using Fibre Channel components and an 8B/10B encoding scheme developed by IBM.
- Broadband Communications proposed a solution using a data scrambling technique.
- Digital Equipment proposed a solution based on research they were doing in gigabit rate networks. It used an 8B/10B encoding scheme with forward error correction (FEC).
- Hewlett-Packard proposed a solution based on research they were doing in gigabit rate data acquisition. It used an 20B/24B encoding scheme.
- TRW proposed a solution also based on data scrambling.

Although all proposed solutions were included in the final voting, there were only three that received support from other members of the group. Those three were the Digital Equipment, Broadband Communications, and Hewlett-Packard proposals. Accordingly, this paper will concern itself only with those proposals.

The Digital Equipment proposal included components from vendors they were working with to develop a complete solution which would provide computer networks that operate at gigabit rates. The proposal required a serial baud rate of 1.5 Gbit/s. It included a forward error correction code that provided for correction of all single bit errors and detection of all errors affecting more than one bit in a frame. The encoding scheme encoded 8 bits of data into 10 bits plus 2 error correction bits.

The Broadband Communications proposal required a serial baud rate of 1.1 Gbit/s. There was no encoding per se. The data bits were delivered in a scrambled order that minimized long runs of ones or zeros. Error detection depended on HIPPI-PH error detection mechanisms.

The Hewlett-Packard proposal also included components from vendors they were working with to develop a complete solution which would provide data acquisition hardware that operated at gigabit rates. The proposal required a serial baud rate of 1.2 Gbit/s. It included a 24 bit encoding of 20 data bits that provided additional error detection beyond that provided by HIPPI.

Both Hewlett-Packard and Digital Equipment described models of errors in serial optical media. These models differed in how errors occurred. The H-P model showed errors occurring in bursts. The DEC model showed errors occurring as uncorrelated single bit errors. Both models were supported by extensive computer modeling and real-world testing experience.

All proposals delivered all the HIPPI data and parity bits, and all HIPPI control signals except for INTERCONNECT which was derived as loss of light.

The Hewlett-Packard solution was selected by the narrowest of margins (one vote) over the Digital Equipment solution. The principal reasons given for the selection were: Hewlett-Packard required less bandwidth (1.2 vs 1.5 Gbit/s). In the experience of several members with optical media, errors occurred in bursts, and burst errors are not correctable by FEC. It was also felt that the FEC approach would be most beneficial in situations where a link was operating at its very maximum length.

5.3: Serial-HIPPI

The Serial-HIPPI can, for the purposes of description, be partitioned into five functional units. The following paragraphs describe each of these functional units.

The HIPPI interface unit converts the 44 HIPPI signal lines (CLOCK, 32 data, 4 parity, REQUEST, PACKET, and BURST in the forward direction, and CONNECT and READY in the reverse direction, plus the 2 INTERCONNECT signals), into 20 bit data fields plus three signals used as input to the link interface.

The link interface encodes the 20 bits from the HIPPI interface into 24 bit frames. The line code provides DC balance, a bit to allow the receiver to differentiate between the first and second 20 bit frames of the encoded 44 HIPPI signals, and a guaranteed transition for clock recovery and frame synchronization.

The link control unit provides the sequence of events needed for link start-up and control. It uses the HIPPI INTERCONNECT signals to report the status of the link to the HIPPI port.

The serial electrical interface provides timing, electrical, and physical requirements for use with coaxial cable.

The serial optical interface provides timing, optical, and physical requirements for use with a single mode fiber.

5.4: Error detection

Error detection is provided for errors occurring in the HIPPI-PH data, as well as for errors introduced by the optical (electrical) media.

Data errors are detected through the use of the HIPPI parity and LLRC. These bits are passed unchanged through Serial-HIPPI, and the errors will therefore be detected by the receiving HIPPI.

Illegal HIPPI source control signal states are detected by the encoding scheme chosen for these signals. All legal REQUEST, PACKET, and BURST states can be encoded in four values. All illegal states encode as "idle". When de-coded this causes the destination HIPPI to receive all these signals as false.

Illegal HIPPI destination control signal states are detected by the HIPPI source as these signals are sent unchanged through Serial-HIPPI.

Loss of an INTERCONNECT signal causes the link control to enter the start-up state. This results in the INTERCONNECT signal being delivered to the other HIPPI being forced false.

Frame errors cause the values of the control lines from the last good frame to be delivered to the receiving HIPPI along with forced parity errors in the data.

The Serial-HIPPI is self-starting, reliable, and will notify the destination of any errors detected in the data or control signals. The probability of erroneous information being delivered as good data is extremely small.

6: Conclusions

HIPPI has been approved as an ANSI standard, and is currently the interface of choice for high-end applications. A distance limitation of 25 meters imposed by the twisted-pair cables has been overcome by specifying a fiber optic based extender called Serial-HIPPI.

7: Acknowledgements

The X3T9.3 Task Group that developed the HIPPI specifications, and the ad hoc group that developed the Serial-HIPPI implementers agreement were made up of excellent people from many different companies. Thirty-nine companies participated in the development of Serial-HIPPI. Don Tolmie of Los Alamos chaired X3T9.3; Marty Halvorson of DEC chaired the Serial-HIPPI ad hoc group. Special mention is given to Bill McFarland of Hewlett-Packard who was the technical editor of the Serial-HIPPI document.

The Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36. This work was performed under auspices of the U.S. Department of Energy.

8: References

[1] HIgh-Performance Parallel Interface - Mechanical, Electrical, and Signalling Protocol Specification (HIPPI-PH) - American National Standard X3.183-1991.

[2] "Serial-HIPPI Specification, Revision 1.0, Serial HIPPI Implementers Group", available via anonymous ftp as file pub/hippi/serial_hippi_1.0.ps on hplsci.hpl.hp.com (15.255.176.57), or as file hippi/serial_hippi_1.0.ps on nsco.network.com (129.191.1.1).

[3] "HIPPI Developments for CERN Experiments," A. van Praag, et al., Presented at IEEE Nuclear Sciences Symposium 1991, Santa Fe, New Mexico.

[4] "Gigabit Network Testbeds," IEEE Computer magazine, September 1990, Vol. 23, No. 9, pp. 77-80.

[5] "Science and Business (Gigabit Connection)," Scientific American, October 1990, pp. 118-120.